Position Statement: Contingent Planning in Partially Observable Stochastic Domains via Stochastic Satisfiability

Stephen M. Majercik[1] and **Michael L. Littman**[2,1]

[1]Department of Computer Science
Duke University
Durham, NC 27708-0129
{majercik,mlittman}@cs.duke.edu

[2]AT&T Labs
180 Park Avenue
Florham Park, NJ 07932
mlittman@research.att.com

Our research has successfully extended the planning-as-satisfiability paradigm to support contingent planning under uncertainty (uncertain initial conditions, probabilistic effects of actions, uncertain state estimation). *Stochastic satisfiability* (SSAT), a type of Boolean satisfiability problem in which some of the variables have probabilities attached to them, forms the basis of this extension. We have developed an SSAT framework, explored the behavior of randomly generated SSAT problems, and developed algorithms for solving SSAT problems (Littman, Majercik, & Pitassi 2000). We have also shown that stochastic satisfiability can model compactly represented artificial intelligence planning domains, an insight that led to the development of ZANDER, an implemented framework for contingent planning under uncertainty using stochastic satisfiability (Majercik & Littman 1999).

ZANDER solves probabilistic propositional planning problems: states are represented as an assignment to a set of Boolean state variables and actions map states to states probabilistically. Problems are expressed using a dynamic-belief-network representation. A subset of the state variables is declared *observable*, meaning that any action can be made contingent on any of these variables. This scheme is sufficiently expressive to allow a domain designer to make a domain fully observable, unobservable, or to have observations depend on actions and states in probabilistic ways. ZANDER operates by solving an SSAT encoding of the planning problem; the solution to this SSAT problem yields a plan that has the highest probability of succeeding. ZANDER can solve arbitrary finite-horizon partially observable Markov decision processes and solves planning problems drawn from the literature at state-of-the-art speeds (Majercik & Littman 1999).

The general motivation for our planning research was to explore the potential for deriving performance gains in probabilistic domains similar to those provided by SATPLAN (Kautz & Selman 1996) in deterministic domains. There are a number of advantages to encoding planning problems as satisfiability problems. First, the expressivity of Boolean satisfiability allows us to construct a very general planning framework. Another advantage echoes the intuition behind reduced instruction set computers; we wish to translate planning problems into satisfiability problems for which we can develop highly optimized solution techniques using a small number of extremely efficient operations. Supporting this goal is the fact that satisfiability is a fundamental problem in computer science and, as such, has been studied intensively. Numerous techniques have been developed to solve satisfiability problems as efficiently as possible. Stochastic satisfiability is less well-studied but many satisfiability techniques carry over to stochastic satisfiability nearly intact (Littman, Majercik, & Pitassi 2000).

There are disadvantages to this approach. Problems that can be compactly expressed in representations used by other planning techniques often suffer a significant blowup in size when encoded as Boolean satisfiability problems, degrading the planner's performance. Automatically producing maximally efficient plan encodings is a difficult problem. This problem has been addressed for deterministic planning domains (Kautz, McAllester, & Selman 1996; Ernst, Millstein, & Weld 1997), but remains unsolved. We are currently exploring the impact of alternative SSAT encodings on ZANDER's efficiency. In addition, translating the planning problem into a satisfiability problem obscures the structure of the problem, making it difficult to use our knowledge of and intuition about the planning process to develop search control heuristics or prune plans. This issue has also been addressed for deterministic domains; Kautz & Selman (1998), for example, report impressive performance gains resulting from the incorporation of domain-specific heuristic axioms in the SAT encodings of deterministic planning problems.

Our current research focuses on two areas: 1) improving ZANDER (better data structures to optimize the application of heuristics, more compact and efficient SSAT encodings, encoding domain knowledge, memoization for contingent planning, using learning to accelerate the solution process, and more sophisticated splitting heuristics), and 2) developing an approximation technique for solving SSAT-encoded planning problems that will allow us to scale up to larger domains.

Acknowledgments: The first author acknowledges the support of a NASA GSRP Fellowship.

References

Ernst, M. D.; Millstein, T. D.; and Weld, D. S. 1997. Automatic SAT-compilation of planning problems. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 1169–1176.

Kautz, H., and Selman, B. 1996. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 1194–1201. AAAI Press/The MIT Press.

Kautz, H., and Selman, B. 1998. The role of domain-specific knowledge in the planning as satisfiability framework. In *Proceedings of the Fourth International Conference on Artificial Intelligence Planning*, 181–189. AAAI Press.

Kautz, H.; McAllester, D.; and Selman, B. 1996. Encoding plans in propositional logic. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR-96)*, 374–384.

Littman, M. L.; Majercik, S. M.; and Pitassi, T. 2000. Stochastic Boolean satisfiability. *Journal of Automated Reasoning*. In press.

Majercik, S. M., and Littman, M. L. 1999. Contingent planning under uncertainty via stochastic satisfiability. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, 549–556. The AAAI Press/The MIT Press.